arXiv:1608.05104v1 [cs.CV] 17 Aug 2016

Scene Labeling Through Knowledge-Based Rules Employing Constrained Integer Linear Programing

Nasim Souly, Mubarak Shah

Center for Research in Computer Vision, University of Central Florida, USA
nsouly@eecs.ucf.edu, shah@crcv.ucf.edu

Abstract

Scene labeling task is to segment the image into meaningful regions and categorize them into classes of objects which comprised the image. Commonly used methods typically find the local features for each segment and label them using classifiers. Afterward, labeling is smoothed in order to make sure that neighboring regions receive similar labels. However, they ignore expressive and non-local dependencies among regions due to expensive training and inference. In this paper, we propose to use high level knowledge regarding rules in the inference to incorporate dependencies among regions in the image to improve scores of classification. Towards this aim, we extract these rules from data and transform them into constraints for Integer Programming to optimize the structured problem of assigning labels to super-pixels (consequently pixels) of an image. In addition, we propose to use soft-constraints in some scenarios, allowing violating the constraint by imposing a penalty, to make the model more flexible. We assessed our approach on three datasets and obtained promising results.

1. Introduction

Semantic segmentation is one of the essential tasks to interpret and analyze an image; objects types and their placement in the scene provide the substantial information about the 3D world captured in the image. The constraints implied by the relations between different type of objects and the geometry of the objects in the scene can aid in processing and understating images. The effectiveness and advantage of using world knowledge in understanding images have been explored and assessed in early vision

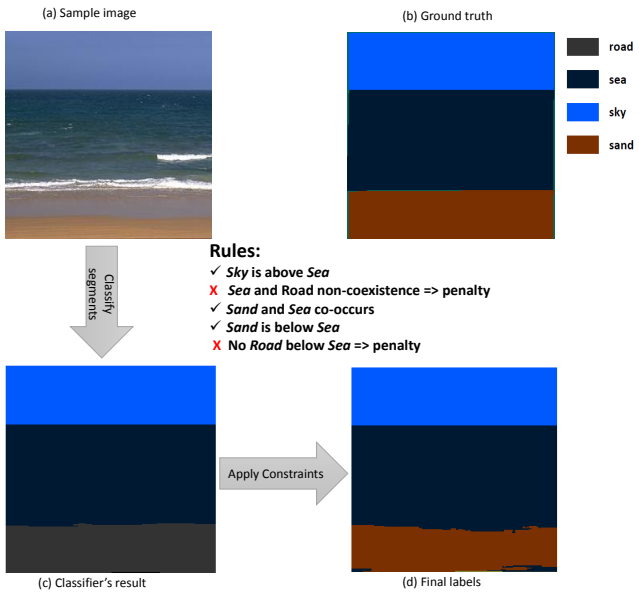

Figure 1: Given an image and initial region labels obtained by the classifiers, we aim to improve the labels using constraints learned from the data. (a) shows a query image, (b) shows the human annotated image (Ground Truth), (c) are labels obtained by classifiers, and (d) shows our results after enforcing constraints.

researches [1], and we want to exploit this concept to label the pixels of an image.

Labeling segments in an image can be considered as assigning labels to a set of variables, which can be pixels or a group of pixels(super-pixels) while considering (or enforcing) the dependency among the labels. Several approaches use linear models for this type of problems. Given features of each super-pixel, for instance as an input, the goal is to use inference to find the best assignments of labels to the super-pixels as output. Many learning methods have been proposed for structured modeling, which mostly attempt to model dependency among labels during learning process by optimizing a global objective function (such as using Markov Random Field and Conditional Random Field formulations). However, due to efficiency and tractability, these methods are confined to encode only local relationships. Although, non-local dependencies can be coded in such a model as well, the model needs to learn more parameters (e.g., in graphical models more edges and weights need to be included to model long-term

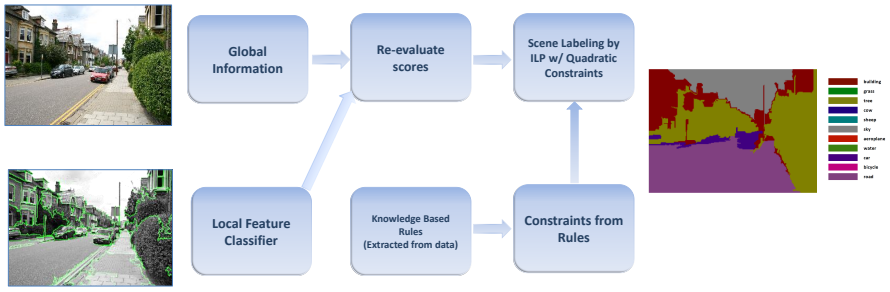

Figure 2: An overview of the proposed approach. For training, we begin by segmenting images into super-pixels and extracting the feature matrix. Then classifiers (extreme gradient boosting trees) are trained. Also we find the scene-labels association matrix to capture global context. In the inference part during testing, for a given image the label scores are obtained via the classifiers results. Finally, labels scores are updated by applying constraints learned from knowledge-based rules through the optimization function by Integer Programming.

dependencies). While this can be done by infusing the knowledge in the model as constraints, rather than using a fully connected graph to capture the all possible interactions or higher order representation can be employed, however, in both cases the complexity of method during learning and inference increases.

In this paper, we propose to use the knowledge of some relevant interactions between labels *directly* in the model, instead of *indirectly* via learning. In doing so, we benefit from the prior knowledge in inference stage, by adding some constraints that should be satisfied. As a result, we apply inference only in the testing, and we do not require solving any inference problem during the training process. Therefore, we can use any training algorithm, and apply constrained inference model; that way the features and constraints can be distinguished. We use a data-driven approach to extract rules to form the training data which is later used to generate expressive constraints. Since the constraints are formed as Boolean functions, they can be represented as logical expressions. By analyzing constraints extracted from the data, we discovered that we need to model soft-constraints in addition to the hard constraints. For instance, most of the time mountain and building are not seen together, but this is not always

the case. Thus, we use slack variables in the objective function to model the soft constraints. To solve the inference problem with expressive constraints, we use an Integer Programming formulation.

By exploiting declarative constraints, we eliminate the need for re-training the model for new incoming data; because we can simply add more constraints in inference part without changing the prediction problem. Furthermore, we exploit global context of the image by learning scene-label association weights, to increase the probabilities of the most confident labels and limit the number of labels, which need to be explored. In order to demonstrate the performance of our method, we report experimental results on three benchmark datasets including SIFTflow [2] MSRC2 [3] and LMSun [4].

In summary, we make following contributions:

- We use extreme gradient boosting algorithm to learn efficient classifiers to label local features with no need to train detectors or expensive classifiers.
- We improve the scores of super-pixels by combining local classifiers results and global scene information obtained by a learned scene-label association.
- We incorporate high-level knowledge of relations between labels via imposing rule-based constraints in an Integer Programming problem.

The rest of the paper organized as follow: section 2 reviews relative work proposed for scene labeling, in section 3 our proposed method is described in details. The experiments and evaluations of our method are presented in section 4, and finally, we conclude our paper.

2. Related Work

Scene labeling has been the interest of many research works recently. Proposed methods differ regarding features and descriptors they use, primitive elements (pixels, patches or regions), employed classifiers and how context techniques are utilized. Conditional Random Fields [5] has been used in a vast number of methods. These works use mainly appearance (local features) as unary potential and smoothness between neighboring elements as the pairwise term [3]. Also using higher order CRF

to integrate potential of features at different levels (pixels and superpixels) have been explored [6], [7]. In addition to local features, some methods benefit from object detectors and combine the results from detectors and context information [5], [4].

Other approaches, labels from a dataset of known labels, a retrieval set, transfer to the segments of the query image. To this aim, for a given image, a nearest neighbor algorithm is used to retrieve similar images from a sample data, then a Markov random field (MRF) model is applied on pixels (or super-pixels) in the image to label the pixels [2] , [8] and [9]. Many works are the extension of this type of labeling, for example, in [10]the weights of descriptors are learned in an off-line manner to reduce the impact of incorrectly retrieved super-pixels. Also, [11] proposed to use a locally adaptive distance metric to find the relevance of features for small patches in the image and to transfer the labels from retrieved candidates to small patches of the image. In [12] instead of using a retrieval set to transfer the labels, a graph of dense overlapping patch correspondences is constructed; and the query image is labeled by using established patch correspondences.

In [13] an image is represented as a collage of warped, layered objects obtained from reference images. The scene is analyzed through synthesizing the information. A dictionary of object segment candidates that match the image is retrieved form samples for a test image; then the image is represented by combining these matches. To do so, a dataset of labels exemplars is required. In [14] authors proposed to find the bounding boxes of the objects using object detectors and then regions are classified by combining information from detectors and surface occlusions; they also use RGB-depth to understand the scene.

In some other works, context information is employed in the modeling of the problem via utilizing global features of the image or incorporating the co-occurrence of the labels in training[15]. Recently, deep learning techniques have been used in scene labeling. For instance, in [16] for each pixel in the image multi-scale features are gained, and the deep network aggregates feature maps and labels the regions with highest scores. These models need an extensive data for training; in scene labeling task pixel-wise annotation is expensive work. Also, unlike expressive models [under-

standing] insight of the relations between labels is hard to achieve [16], [17].

Most of these approaches neglect to use domain knowledge and the structure -the interactions and dependencies between different parts of the image geometrically (spatially) or conceptually- implied by the images, or they only consider the smoothness of neighboring pixels. While the early researches in computer vision (e.g., VISIONS [1]) tried to interpret scenes using knowledge from the structure of the image. In this paper, we aim to incorporate modeling constraints into inference to leverage from this type of information.

Previous works mainly use fully connected graphs to model long-distance dependencies in contrast to local potentials such as MRF models, in which pairwise terms usually are based on the first-order Markov assumption due to reducing the complexity of the inference. However, increasing the number of edges between nodes in a graph causes the inference to be competently expensive. Our goal is to use knowledge-based information to discard the unnecessary edges and make the model simple. We aim to capture the real and common dependencies using the rules extracted from the data instead of letting the CRF finds the weights for these dependencies. Also, our method can be general in a sense that the various types of relations can be formulated in the same framework of rules such as geometrical dependency, non-co-existence, co-occurrence, presence, and adjacency. While in other works (e.g. [18], [19]. and [20]) only a subset of these constraints are modeled as hard constraints, whereas the rules which extracted from the real data show that some of the constraints are not hard constraints. Therefore, we use soft constraints for those rules that are most of the time valid, but not always.

3. Approach

Our system consists of two main phases. The first phase consists of off-the-shelf parts including feature extraction and classifier training based on local features of the sample training images. The second phase is the inference, in which for a given query image, using scores computed by the classifiers for each possible label, an objective function is maximized such that the constraints are not violated. An overview of our

proposed approach is shown in figure 2.

In training, first we segment images using efficient graph-based segmentation [21]. We use this specific method to be consistent with other methods we compare our method with. Next, for each super-pixel, local features, including SIFT, color histogram, mean and standard derivation of color, area and texture, are extracted. Given these local features, classifiers (extreme gradient boosting trees forest) are trained to label super-pixels using their local features.

In the inference part, we update the classifiers scores by reducing the scores of the assignments which conflict with the defined constraints. We use the product of experts ([22], [23]) by combining the probabilities of the label assignments given by the classifiers and the degree of inconsistency provided by the constraints. This can be done by multiplying the probabilities or equivalently adding their logarithm. Note that this is different from using various features and weighting those features. We define a new scoring function:

$$score(\mathrm{y}^i, \mathrm{x}^j) = \Phi(\mathrm{y}^i|\mathrm{x}^j) \times C(\mathrm{y}^i), \tag{1}$$

where Φ is classifier probabilities and C is constraints violation degrees.

We formulate the prediction of a label for each segment in the image by constrained optimization problem, and employ Integer Linear Programing (ILP) with quadratic constraints as an optimization framework. In doing so, we assign a binary label $y_i^j \in \{0, 1\}$ to a super-pixel i, in order to find the best assignment $\mathcal{Y} = \{y_1^j, y_2^j, ..., y_n^j\}$, where n is the number of super-pixels and j belongs to labels $1, ..., l$. Hence, given the classifier score for each super-pixel, we formulate the inference function as

$$\arg\max_{\mathrm{y}} \sum_{i,j} \mathrm{w}(i, j)\ y_i^j, \tag{2}$$

where $\mathrm{w}(i, j)$ is the score of assigning label j to super-pixel i, provided by our local feature based classifier. This objective function maximize the number of correct predictions for the super-pixels in the image. Toward exploiting context and prior knowledge,

we introduce different types of constraints such as non-coexistence, presence and etc, which are explained in 3.3 . The first constraint, which should be applied to all instances, enforces that each super pixel gets exactly one label.

$$\forall i \ \sum_{j=1}^{n} y_i^j = 1. \tag{3}$$

In following sections, we explain each part of the approach in detail.

3.1. Local Classifiers

In this section, we explain the first step of our method. In training, we start with segmenting each sample image into super-pixels using efficient graph-based segmentation method [21], followed by computing a feature vector(including, SIFT, color mean) for each super-pixel in the image. We use the same features as used in [8].

We use a sigmoid function to rescale the classifier scores to give a chance to other classes, beside the one the classifier with a maximum score, to compete during the optimization phase. By doing so, if the classifier mislabels a super-pixel there is a chance that the label may be changed during the inference phase by applying the constraints. We adapt the parameters of the function using the validation sample data. Also, the sizes (areas) of super pixels are multiplied by the scores to make larger super-pixels more important during the optimization.

We use Extreme Gradient Boosting [24] with softmax objective function to categorize each super-pixel in an image. Since the training data inevitably is noisy, super-pixels may break the structure of the data, the bagging using a subset of training examples and subsets of features are used to reduce the effects of the noisy data. Unlike some of the other methods, which train object detectors in addition to the region classifiers, we only use region features and simple classifiers to obtain the initial label scores for super-pixels. In our experiments, boosting trees achieved better results in terms of average accuracy among all the classes, even though we discard some of the samples randomly during the training. We discard some samples since the number of sample data for some classes is enormous.

3.2. Global Context Information

In performing scene labeling, incorporating scene information, such as the places, can be helpful. For instance, *if the probability of being a desert is high for an image presumably the chance of seeing chair or river would be low*. Therefore, we use image level categories to refine the scores of the local classifiers. We use Places CNN model [25] to find the most probable scene semantics in a given image. Unlike pixel-level annotation, image level annotation is more feasible, thus training a deep network is doable. This deep network is trained on scene labeled images and includes 205 places categories. Using our training data, we learn a mapping between these categories and the label set. To do so, we employ non-negative sparse regression formulation, and extract a weight matrix W for scene-label association as follow:

$$\min_{W \geq 0} \|Y - WX\|_F^2 + \lambda \|W\|_1, \tag{4}$$

where X is the confidence score matrix for scene-categories of training images, and Y is a matrix of present labels in corresponding images. In the interest of putting emphasis on learning the mapping between scene categories and smaller super-pixels and rare classes, in Y if label l_i is present we use $1 - n(l_i)/\sum_j n(l_j)$ in i_{th} element. The solution of this problem can be efficiently obtained by FISTA algorithm [26] which is implemented in SPArse Modeling Software(SPAMS) [27].

Then in testing time, for a given image, using scene categories and the weight matrix the most confident labels are obtained and used to refine the label assignment.

3.3. Extracting Rules and Creating Constraints

In this section, we describe the types of constraints that we add to the aforementioned optimization problem. Note that even though these constraints can be obtained by common sense knowledge, we explore the training data to discover plausible rules in the context of semantic labeling of images. For instance, the common sense about label *sky* and *building* would be: *sky always is above the building*; however, since images are 2D projections of a 3D world this is not always correct. Therefore, we capture these types of constraints as relative constraints with some penalty if they are violated.

We call them soft-constraints and describe them in section 3.4

The first type of constraints are **geometrical constraints**; whether two labels can have a particular geometrical relationship or not. For example, *sky* always is seen above the *sea*. For each super-pixel, we keep the bounding boxes information, then we apply a rule: *When the label for a super-pixel is sea, its lower super-pixels can not get label sky*. This type of constraints can be formulated as follows:

$$y_i^a \sum_{j=i+d}^{n} y_j^b \leqslant 0, \tag{5}$$

where a and b respectively are labels of the lower and upper parts of the image, and d is the displacement of super-pixel j from super-pixel i. By exploring the data, we can find hard constraints belonging to this group (e.g (sky, field), (sea, sand), (ceiling, floor)...).

The second type of constraints are **mutually-exclusive**, which represent cases *when labels a and b, never occur together in an image*. This can be formulated as follows

$$\sum_i y_i^a \sum_j y_j^b \leqslant 0. \tag{6}$$

Presence is another type of constraints which can be applied. According to this constraint, *if a certain label a appears in the image, then there should be at least one super-pixel having label b.* We can express this constraints as one of the following formulations:

$$\sum_i y_i^a \sum_i y_i^b \geqslant 1, \tag{7}$$

$$\sum_i y_i^a - \sum_{i,j} y_i^a y_j^b \leqslant 0 \tag{8}$$

For example, we have found that whenever the label *balcony* appears in the image there exists, at least, one region labeled as a *building*, similarly whenever the label *pole* appears in the image, at least, one region is labeled as a *road*. Note that; this constraint is different from "co-occurrence. In co-occurrence two labels frequently appear together in images; the constraint is not necessarily always valid, thus, it can suitably

be formulated as a soft-constraint.

These rules can be extended, for example, we could add adjacency constraints to enforce the neighboring super-pixels receive the same labels, or even we can apply some constraints on the other features of the specific labels.

We define the rules as described in this part, and then we use a sample data to find relations between classes which follow these rules. We create a cube which with the size of $numberoflabels \times numberoflabels \times numberofrelations$. Each matrix of this cube shows the frequency of a relation happening between classes, and each cell indicates the frequency of a particular relationship between two categories (class labels). For example, to find labels that have a specific relation, we count how many times this relation occurs in a sample data for each pair of classes. If the value is higher than a threshold, we will consider it as a constraint. In addition, we distinguish between soft and hard constraints if the other way around has occurred or not. For instance, about the geometry relation like above, for a pair semantic labels (ceiling, floor) the relations is always true. Therefore, we consider it as hard constraints. While, for (sky, building) there are some images in which the building super-pixels are above sky, thus we consider this relation as soft constraint, it means that the main formulation can tolerate violating this constraint with a penalty.

3.4. Integer Linear Programming with soft constraints

Our proposed approach for extraction of rules not only helps us in deriving the hard constraints, but it also contributes to providing the soft constraints such as *co-occurrence* and *relative geometrical* constraints.

Since some of the rules may not necessarily be satisfied by all the data, we use 0-1 soft-constraint modeling [28] and define soft constraints by introducing a new binary variable, z_k, for each constraint and an associated penalty, c, which indicates the degree of violation (how the confidence of assignment should reduce when the constraint is not satisfied). Then objective function in 2 becomes:

$$\arg\max_{y} \sum_{i,j} \mathrm{w}(i,j)\; y_i^j - \sum_{k} c_k(1 - z_k), \tag{9}$$

which implies that if constraint k is violated z_k would be zero and consequently a penalty will be imposed to the optimization function. Moreover, we need to connect the z_k to the constraint C_k as $z_k \leftrightarrow C_k$. We are able to do that by using logical representation and adding these constraints into the objective function. For example, *sky mostly is above the building*; however, based on the rules that we have extracted from the database, this is not always true. Therefore, we change the constraint to soft constraint as

$$y_i^{building} \sum_{j=i+d}^{n} y_j^{sky} = 0 \leftrightarrow z_k, \tag{10}$$

here $\leftrightarrow$ is an equivalence for the conditional constraints, meaning that if constraint k is violated, z_k will be zero. Consequently, in the equation 9 $c_k(1 - z_k)$ will be a positive number, which is subtracted from the score. In order to formulate if-else condition in the Integer programming we use the binary variable and add the constraints as inequalities, which is a common practice.

The penalty values are obtained statistically from the dataset by finding the probability of each case divided by the number of their appearances. For example, *sky* and *building* most of the time occur together, but not all the time. Therefore, we find a penalty using the following formula, which takes into account the frequency of the constraint violations in the data,

$$c_k = -\log \frac{P(C_k = 0)}{P(C_k = 1)}, \tag{11}$$

where $P(C_k = 0)$ and $P(C_k = 1)$ are respectively the probability of violating and satisfying the constraint k.

3.5. Solving Integer Programming

Thanks to many available numerical packages for estimating Integer Programming solutions, we are able to solve the problem of our size in a short time. We use Gurobi toolkit [29], which can solve 70 IQP per second on a desktop computer. The solver

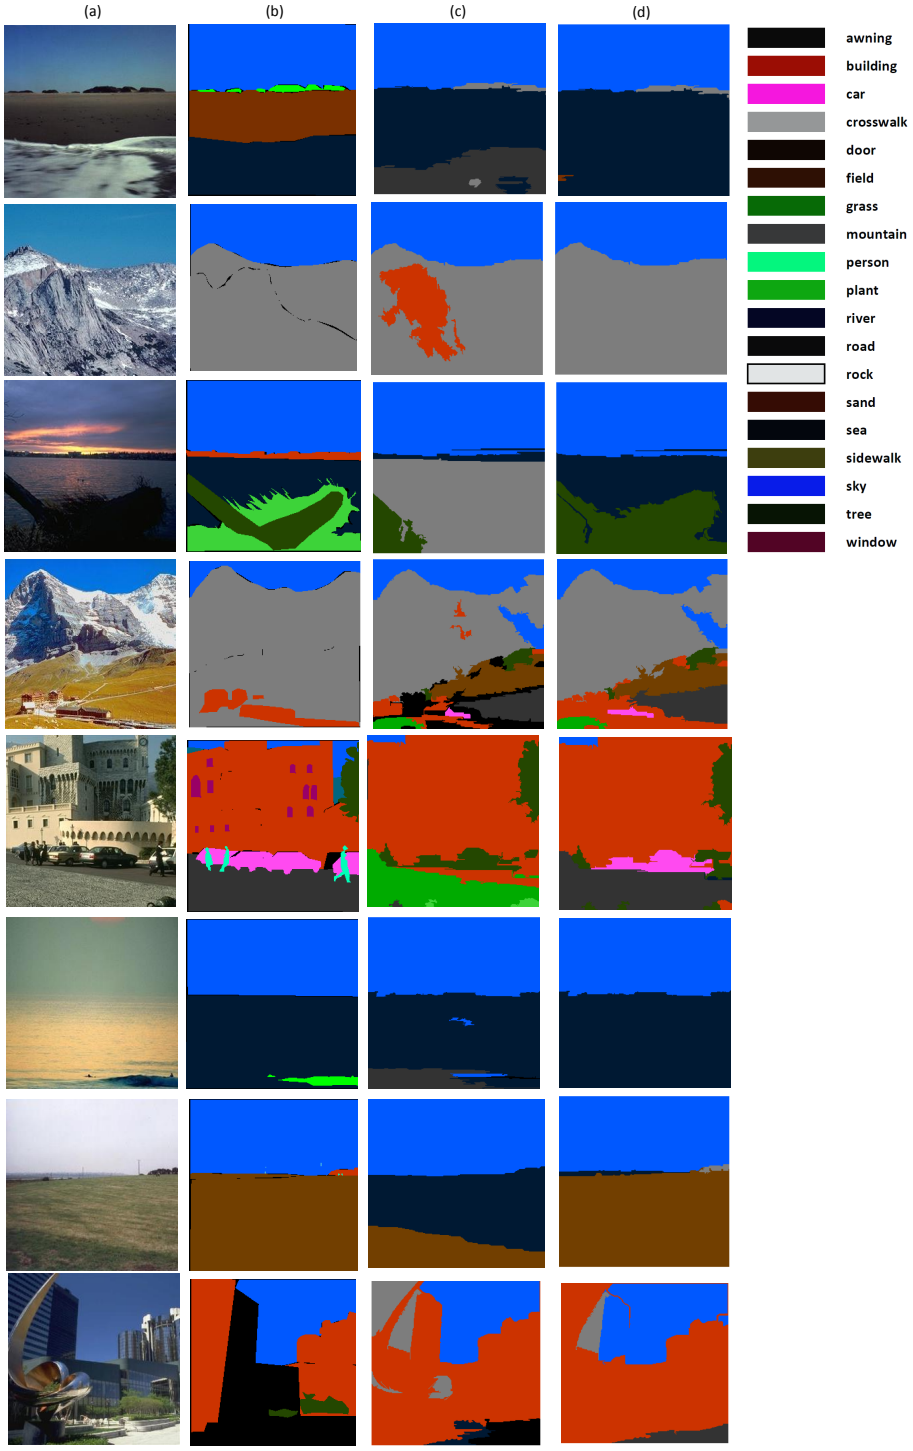

Figure 3: Examples results obtained by our method on SIFTFlow dataset, (a) query images, (b) ground truths, (c) initial classifier outputs and (d) our final results.

Table 1: **Summery of the rules extracted from the sample data.**

Name	Logic representation	examples
non-coexistence	$\mathbf{y}^{l_1} \wedge \mathbf{y}^{l_2} = 0,\ \mathbf{y}^k = (y_1^k \vee y_2^k \vee ... \vee y_n^k)$	(desert, sidewalk)
geometric	$y_i^{l_1} \wedge \mathbf{y}_{i+z}^{l_2} = 0,\ \mathbf{y}^k = (y_{i+z}^k \vee ... \vee y_n^k)$	(road below window)
presence	$y_i^{l_1} \Rightarrow \mathbf{y}^{l_2},\ \mathbf{y}^k = (y_1^k \vee y_2^k \vee ... \vee y_n^k)$	(balcony, building)
co-occurrence	$\mathbf{y}^{l_1} \wedge \mathbf{y}^{l_2} = 1,\ \mathbf{y}^k = (y_1^k \vee y_2^k \vee ... \vee y_n^k)$	(car, road)
adjacency	$y_i^{l_1} \equiv y_j^{l_2} = \neg(y_i^{l_1} \oplus y_i^{l_2}) = 1$	(sky, sun)

uses a piece-wise linear optimization and relaxed LP to solve the integer programming. Also, it is feasible to convert the quadratic constraints to linear ones by adding slack variables due to the boolean nature of our problem. Moreover, one may use other greedy or heuristic methods to solve the objective function.

4. Experiments and Results

We use the SIFTFlow dataset [2], which consists of 2,488 train and 200 test images collected from LabelMe[30]. These images are from the broad categories such as natural images (coasts, country, etc.) and urban images. Ground truth images are available, in which each pixel of the image is labeled from 33 classes including different objects (e.g., cars, windows) and stuff (e.g., sky, sea). In table 3, the accuracy results of our method compared to the state of the similar methods (regarding using similar features) are shown. Our method obtains comparable per-pixel accuracy; even though we do not use object detectors or massively training algorithms. Also in table 2, we show the accuracy gained in different steps of our method, indicating the improvement due to addition of constraints.

In figure 3, we show some qualitative results from SIFTflow data set including sample images, ground truth labeling, results from [4], and results obtained by our method before optimization and after applying constraints based optimization. As it is clear, we are able to improve the results without over smoothing. Also, our method achieves promising results in terms of per-class accuracy as well. For instance, in the

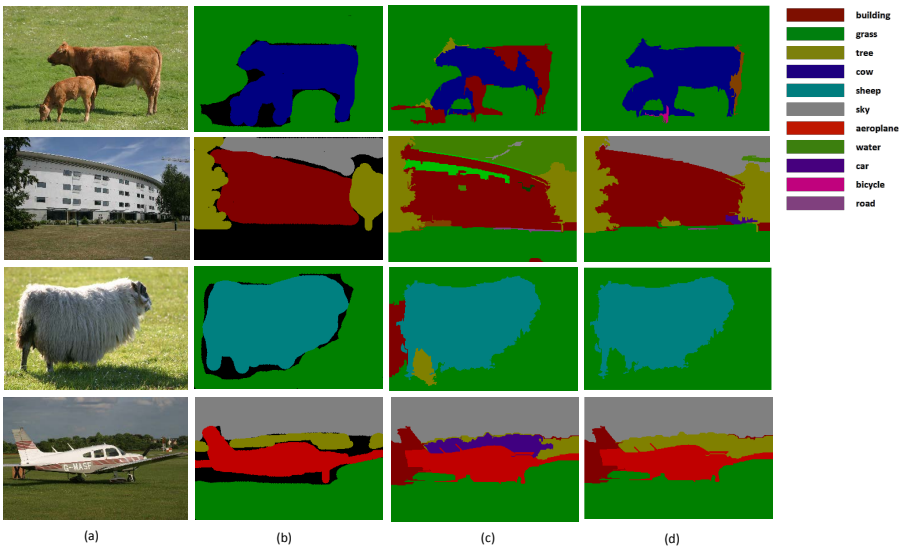

Figure 4: Example results obtained by our method on MSRC dataset, (a) query images, (b) ground truths, (c) initial classifier output and (d) our final results.

third-row super-pixels below the sea are labeled as mountain by our local classifiers, yet constraints such as *mountains are not below sea*, are able to handle the miss-labeling and change it to the *tree*. Note that plant and tree are very similar in terms of the appearance.

Table 2: Detailed Results on SIFTFlow dataset

Method	Avg Accuracy
Local Features + Classifiers	72.8
Local Features + Global Context	77.3
Local + Global + Rules	80.9

The second data set which we assessed our approach on is LMSun [4]. This data set contains 45,676 training images and 500 test images including indoor and outdoor scenes, with different image sizes. However, the number of available samples for different labels are not balanced, and some labels are rare. While some labels, for instance, sky and building, have a large number of samples, some others have fewer

Table 3: Accuracy on SIFTflow dataset

Method	Per-Pixel	Per-Class
Farabet natural [16]	78.5	29.6
Farabet balanced [16]	74.2	46.0
Tighe [4]	78.6	39.2
Collage Parsing [31]	77.1	41.1
Gerorge w/out Fisher Vectors[32]	77.5	47.0
Gerorge Full [32]	81.7	50.1
Ours	80.9	50.3

samples. To make the training more feasible, we use all samples from rare classes and only 25% of samples from common classes for training our extreme gradient boosting classifiers. Also, in learning as well as while using the scene-label associations we assign more weights to smaller super-pixes and rare classes to avoid the influence of common labels such as sky or building.

Table 4: Accuracy on LMSun dataset

Method	Per-Pixel	Per-Class
Tighe [4]	61.4	15.2
Gerorge w/out Fisher Vectors[32]	58.2	13.6
Gerorge Full [32]	61.2	16.0
Ours Local Classifiers	47.4	13.4
Ours Local Classifiers +Global Context	56.1	14.6
Ours	58.4	17.3

We also applied our method on MSRCV2 data set [3] which has 591 images of 23 classes. We use the provided split, 276 images in training and 255 images for testing. The qualitative results are shown in figure 4 and our result using the aforementioned approach is presented in table 5.

Table 5: Detailed results for MSRCV2 dataset

Method	Per Pixel Accuracy
Local Features + Classifiers	76.7
Local Features + Global Context	78.3
Local + Global + Rules	84.4

5. Discussion

The proposed method can get similar results compared to other models such as CRF by having fewer edges i.e. not considering fully connected graph. The two sources of information, including visual features and high-level knowledge-based constraints, are combined to obtain better results. Some relevant information about the data, such as geometrical relationships or non-coexistence, which is hard to learn from the features automatically, can be easily captured by our proposed method and used to refine the labels. Our approach can scale and generalize.

In our approach, we add constraints to the trained model, so without retraining, we can add constraints which are declarative and hard to model using only features. Hence, in cases when not enough training data for some categories (class labels) is available, the constraints will help us to obtain better results. Also, when new classes are added to the data set, the proposed method can model the dependencies between classes without learning the pairwise terms since we keep the features (classifier learning) and constraints separate. As shown in experiments on different data sets, with a various number of labels, our method can get promising results. The constraints are assumed to be trusted, and penalties are obtained from data by simply finding the frequency in contrast to learning from the features. It should be noted that our primary contribution is improving the labeling on top of classifier scores; therefore, using extensive classifiers, such as deep learning, can boost the final results.

6. Conclusions

In this paper, we proposed a novel scene labeling approach, in which we use an enhanced inference method that enables the model to incorporate general constraints

structure. We use an integer programming formulation, which can be solved by linear programming relaxation to address the problem. We also proposed to use soft constraints in addition to hard constraints to make the model more flexible.Experimental results on three data sets show the effectiveness of our method.

References

References

[1] C. C. Parma, A. R. Hanson, E. M. Riseman, Experiments in schema-driven interpretation of a natural scene, Springer, 1981.

[2] C. Liu, J. Yuen, A. Torralba, Sift flow: Dense correspondence across scenes and its applications, Pattern Analysis and Machine Intelligence, IEEE Transactions on 33 (5) (2011) 978–994.

[3] J. Shotton, J. Winn, C. Rother, A. Criminisi, Textonboost for image understanding: Multi-class object recognition and segmentation by jointly modeling texture, layout, and context, International Journal of Computer Vision 81 (1) (2009) 2–23.

[4] J. Tighe, S. Lazebnik, Finding things: Image parsing with regions and per-exemplar detectors, in: Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on, IEEE, 2013, pp. 3001–3008.

[5] L. Ladickỳ, P. Sturgess, K. Alahari, C. Russell, P. H. Torr, What, where and how many? combining object detectors and crfs, in: Computer Vision–ECCV 2010, Springer, 2010, pp. 424–437.

[6] C. Russell, P. H. Torr, P. Kohli, Associative hierarchical crfs for object class image segmentation, in: in Proc. ICCV, Citeseer, 2009.

[7] S. Gould, R. Fulton, D. Koller, Decomposing a scene into geometric and semantically consistent regions, in: Computer Vision, 2009 IEEE 12th International Conference on, IEEE, 2009, pp. 1–8.

[8] J. Tighe, S. Lazebnik, Superparsing: scalable nonparametric image parsing with superpixels, in: Computer Vision–ECCV 2010, Springer, 2010, pp. 352–365.

[9] J. Tighe, M. Niethammer, S. Lazebnik, Scene parsing with object instances and occlusion ordering, in: Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on, IEEE, 2014, pp. 3748–3755.

[10] D. Eigen, R. Fergus, Nonparametric image parsing using adaptive neighbor sets, in: Computer vision and pattern recognition (CVPR), 2012 IEEE Conference on, IEEE, 2012, pp. 2799–2806.

[11] G. Singh, J. Kosecka, Nonparametric scene parsing with adaptive feature relevance and semantic context, in: Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on, IEEE, 2013, pp. 3151–3157.

[12] S. Gould, Y. Zhang, Patchmatchgraph: Building a graph of dense patch correspondences for label transfer, in: Computer Vision–ECCV 2012, Springer, 2012, pp. 439–452.

[13] P. Isola, C. Liu, Scene collaging: Analysis and synthesis of natural images with semantic layers, in: Computer Vision (ICCV), 2013 IEEE International Conference on, IEEE, 2013, pp. 3048–3055.

[14] R. Guo, D. Hoiem, Labeling complete surfaces in scene understanding, International Journal of Computer Vision (2014) 1–16.

[15] T. L. Vu, S.-W. Choi, C. H. Lee, Improving accuracy for image parsing using spatial context and mutual information, in: Neural Information Processing, Springer, 2013, pp. 176–183.

[16] C. Farabet, C. Couprie, L. Najman, Y. LeCun, Scene parsing with multiscale feature learning, purity trees, and optimal covers, arXiv preprint arXiv:1202.2160.

[17] R. Socher, C. C. Lin, C. Manning, A. Y. Ng, Parsing natural scenes and natural language with recursive neural networks, in: Proceedings of the 28th international conference on machine learning (ICML-11), 2011, pp. 129–136.

[18] L. Ladicky, C. Russell, P. Kohli, P. H. Torr, Graph cut based inference with co-occurrence statistics, in: Computer Vision–ECCV 2010, Springer, 2010, pp. 239–253.

[19] A. Roy, S. Todorovic, Scene labeling using beam search under mutex constraints, in: Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on, IEEE, 2014, pp. 1178–1185.

[20] A. Rabinovich, A. Vedaldi, C. Galleguillos, E. Wiewiora, S. Belongie, Objects in context, in: Computer vision, 2007. ICCV 2007. IEEE 11th international conference on, IEEE, 2007, pp. 1–8.

[21] P. F. Felzenszwalb, D. P. Huttenlocher, Efficient graph-based image segmentation, International Journal of Computer Vision 59 (2) (2004) 167–181.

[22] G. E. Hinton, Products of experts, in: Artificial Neural Networks, 1999. ICANN 99. Ninth International Conference on (Conf. Publ. No. 470), Vol. 1, IET, 1999, pp. 1–6.

[23] M.-W. Chang, L. Ratinov, D. Roth, Structured learning with constrained conditional models, Machine learning 88 (3) (2012) 399–431.

[24] J. H. Friedman, Greedy function approximation: a gradient boosting machine, Annals of statistics (2001) 1189–1232.

[25] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, A. Oliva, Learning deep features for scene recognition using places database, in: Advances in neural information processing systems, 2014, pp. 487–495.

[26] A. Beck, M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, SIAM journal on imaging sciences 2 (1) (2009) 183–202.

[27] J. Mairal, F. Bach, J. Ponce, Sparse modeling for image and vision processing, arXiv preprint arXiv:1411.3230.

[28] V. Srikumar, Soft constraints in integer linear programs.

[29] I. Gurobi Optimization, Gurobi optimizer reference manual (2015).
URL http://www.gurobi.com

[30] B. C. Russell, A. Torralba, K. P. Murphy, W. T. Freeman, Labelme: a database and web-based tool for image annotation, International journal of computer vision 77 (1-3) (2008) 157–173.

[31] F. Tung, J. J. Little, Collageparsing: Nonparametric scene parsing by adaptive overlapping windows, in: Computer Vision–ECCV 2014, Springer, 2014, pp. 511–525.

[32] M. George, Image parsing with a wide range of classes and scene-level context, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 3622–3630.